Statistical Modeling for Radiation Hardness Assurance: Toward Bigger Data

R. Ladbury and M. J. Campola

Abstract— New approaches to statistical modeling in radiation hardness assurance are discussed. These approaches yield quantitative bounds on flight-part radiation performance even in the absence of conventional data sources. This allows the analyst to bound radiation risk at all stages and for all decisions in the RHA process. It also allows optimization of RHA procedures for the project's risk tolerance.

Index Terms—probabilistic risk assessment, radiation effects, reliability estimation, quality assurance, radiation hardness assurance methodology

I. INTRODUCTION

Because radiation testing is destructive, statistical models are needed to bound flight part performance using test data for a sample representative of flight parts. Such models can bound failure/error rates, degradation or performance anomalies, all of which are needed for reliability estimates, part selection, design and other Radiation Hardness Assurance (RHA) activities.

Historically, RHA statistical models could use only the most representative data. For Total Ionizing Dose (TID) and Displacement Damage Dose degradation, this meant data for the flight wafer diffusion lot. For Single-Event Effects (SEE), this usually meant data for a sample made with the same mask set and fabrication process as flight parts (not necessarily the same lot}. Other types of data, such as TID historical data (same part type) or data for similar part types manufactured in the same process (so-called similarity data, used for SEE or TID) served as *qualitative* guides for flight part performance. One reason for this is that disentangling the contributions of part-to-part, lot-to-lot and part-type-to-part-type variability when using historical and similarity data requires a complicated statistical model.

R. Ladbury is with NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA (phone: 301-286-1030; fax: 301-286-4699; e-mail: Raymond.L.Ladbury@nasa.gov).

Michael J. Campola is with NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA, Micael.J.Campola@nasa.gov).

Here we discuss several statistical techniques and models that have proven useful in RHA efforts. The methods fall into two categories. We call methods using only the most representative data conventional methods. Although the statistics for such conventional models are well understood, we have developed techniques that elucidate the sensitivities of RHA results to errors in the data or model. For SEE, we discuss techniques for fitting test data and bounding SEE rates. For TID, we consider how to identify "problem parts," where part radiation response does not follow a well behaved distribution. These techniques facilitate test planning and execution and allow the analyst to verify that the data conform to the assumed statistical model. Although a survey of past and current approaches to TID and SEE RHA is beyond the scope of this work, such reviews exist in the literature.[4,5]

Unfortunately, there are some things conventional methods cannot do. They can only bound flight part performance (e.g. for use in design and part selection) once lot-specific TID or part specific SEE test data become available. They do not provide a context for interpreting whether a test result is expected or not. To address these issues, we are developing Bayesian approaches to RHA, which can use many types of data in a single coherent framework to bound radiation performance of candidate parts at all stages of the design process. We consider flight-lot data, historical data, similarity data, and heritage data—all of which can also provide a context for interpreting test results as they are realized. The techniques can be generalized to other types of data as well.

Using larger datasets poses challenges. For example, data may be harvested from multiple sources, some perhaps less controlled than ideal. Flexible statistics must be developed that apply for all the data and still bound flight-part radiation performance. A flexible statistic allows incorporation of more data—which is important, since RHA often relies on small samples and sparse data. It is also important to understand the question(s) the data are answering. For instance, even flight lot test data can give only a conditional answer as

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

1

to how the flight parts will perform: Applying one-sided tolerance limits to find the 99% worst-case (WC) radiation performance (RP) of a part from our flight lot with 90% confidence (RP(99/90)), really says: If the test sample used is in family with at least 90% of possible similar samples (the confidence level—CL—in our data), and as long as the flight parts perform at least as well as 99% of parts (the probability of success, P_s) from the resulting distribution fit to our data sample, the RP of the flight parts should not be worse than RP(99/90). We will briefly summarize what sorts of questions can be answered with historical, similarity, heritage and other data.

Lastly, we look at statistical techniques that can tell us when data deviates from statistical model assumptions in ways that invalidate the analysis and consider how to present results for clarity and relevance.

II. CONVENTIONAL SEE HARDNESS ASSURANCE

In conventional RHA, variation in part-to-part and lot-to-lot SEE response is usually treated as negligible—either because the test data are for the flight lot, or because process variations are assumed to have a limited effect on the SEE rate.   The dominant errors in the SEE cross section (σ) vs. Linear Energy Transfer (LET) curve are Poisson fluctuations in event counts for each LET, or equivalently, fluctuations in failure fluence for destructive SEE. (Ref. [1] treats the case where fluence errors at each LET are also significant.) The σ vs. LET curve is the device input to the rate estimation tools. Usually, this input involves fitting the σ vs. LET data to a family of curves—e.g. cumulative Weibull or Lognormal forms—and the fit parameters serve as inputs to the rate estimation tool.

The usual guidance for dealing with Poisson errors on cross-section measurements is to accumulate sufficiently large event counts that the errors are negligible.[2,3] Then one can fit the cross-section data without reference to the errors.   Unfortunately, many common fitting techniques (e.g. ordinary least squares, or OLS) do a poor job of fitting σ vs. LET, especially near onset, because σ varies by several orders of magnitude and the data points near saturation dominate most goodness of fit (GOF) metrics. This has led to many strategies for fitting SEE data, ranging from "by eye" to weighted OLS. Empirically, we have found that the metric in (1), which we refer to as least log squares (LLS) does a better job of finding correct onset LETs (see Fig. 1) than OLS, and it also does well finding the other parameters:

$$LLS = \sum_{i=1}^{k} \sigma_i^{exp}\left(\ln(\sigma_i^{exp}) - \ln(\sigma_i^{obs})\right)^2 \qquad (1)$$

In (1), $\{\sigma_i^{obs}\}$ is the set of observed cross sections—one for each of k LET values—and $\{\sigma_i^{exp}\}$ are the expected cross section from our model—e.g. a Weibull form. LLS is related to the G statistic and the Kullback-Liebler divergence between two distributions.[6]


Fig. 1 The LLS goodness of fit metric outperforms OLS for determining fit parameters that depend on the near threshold behavior of the fit function—as shown here for onset LET for a Weibull fit to Monte Carlo generated data.

Unfortunately, large event counts are not always possible. Destructive SEE data may require one part per event, making large counts impractical.   Some SEE modes may be so disruptive or rare that they would require very long test campaigns. In some cases, TID degradation or latent damage may preclude gathering large statistics. Then Poisson errors cannot be ignored.

Reference [7] presented a flexible method for treating Poisson errors in SEE σ vs. LET fits using generalized linear models (GLM). In a GLM, variability of a statistic x (e.g. event counts) about its mean μ is described by a member of the exponential family of distributions $P(x,\mu)$ (e.g. the Poisson distribution). The mean μ is described by the model to be fit to the data. The best-fit model parameters, $p_1, p_2...p_k$ are then determined by maximizing the likelihood, $\mathscr{L}$

$$\mathscr{L} = \prod_{i=1}^{n} P(x_i, \mu(p_1, p_2...p_k)) \qquad (2)$$

as a function of the model parameters, for example, σ_{lim}—the limiting cross section, LET_0—the onset LET, and Weibull parameters width W and shape s.

Because $\mathscr{L}$ decreases normally moving away from best-fit values, we establish confidence intervals for parameters using the inverse χ^2 distribution with degrees of freedom equal to the number of model parameters:

$$\ln(\mathcal{L}(CL))/\mathcal{L}_{Max} = -0.5 * INV\chi^2(1-CL,\#param) \quad (3)$$

where $\mathcal{L}_{Max}$ is the maximum likelihood and $\mathcal{L}(CL)$ is the value of $\mathcal{L}$ that defines the CL in the parameter space.

Using (3), one can determine the most likely rate as well as the worst-case SEE rate consistent with all possible fits to the data for a desired confidence. References [7] and [8] illustrated use of the method for test planning. A good metric for assessing a data set's statistical errors is the ratio of the 90%WC SEE rate to the best-fit rate. Usually, the strategy that reduces this ratio the most leads to the best dataset for hardness assurance. Fig. 2 illustrates this process using single-event latchup (SEL) data for the Linear Technology Corp. (LTC) LTC1419 Analog-to-Digital Converter (ADC). The figure shows SEL σ vs. LET data with error bars and the best and 90% worst-case fits to the data determined via the GLM technique outlined above. The table below the figure shows the ratio of the rates resulting from the 90% WC fit to the best fit if we add no data (column labeled Current), if we add data at low LET on the low LET cross section and better determine LET_0, at high LET to better determine σ_{lim} or both low LET and high LET. (Note: for this analysis, we have not assumed any new ions or angles were used, but instead added data to existing cross section estimates to reduce error bars.)

Faster computer speeds have made it practical to use this method in real time for test decisions on the fly during a test run (e.g. whether to test at high LET to better establish σ_{lim} or low LET to establish LET_0). We have implemented the GLM as a spreadsheet that can run in Microsoft Excel™ on a laptop computer.

III. DESTRUCTIVE SEE

Destructive SEE (DSEE) merit special consideration. Because of their severe consequences and the difficulty of estimating accurate failure rates for destructive SEE, the most common risk mitigation approach for these modes has been risk avoidance. This has meant rejecting susceptible parts (e.g. for SEL) or using them only under conditions where susceptibility is negligible (e.g. within the safe operating area for single-event gate rupture—SEGR—or single-event burnout—SEB—for power MOSFETs). Unfortunately, this is not possible in all applications. Some components afford performance advantages that must be weighed against SEE-induced failure risks. Sometimes—e.g. for legacy hardware—a component's susceptibility may come to light after it is designed into the hardware. Required deratings may compromise performance unacceptably. These factors

have increased interest in reliable destructive SEE rate estimation.[9-13] Such methods usually require estimation of at least the σ_{lim} for failure or of σ vs. LET.


Current	Better LET$_0$	Better σ_{lim}	Both
4.5	3.6	3.3	2

Fig. 2 Because the onset LET for SEL in the LTC1419 ADC is between 55 and 59 MeVcm^2/mg, the σ_{SEL} vs. LET curve has only 3 data points. The figure shows the best and 90% WC fits to the data, while the table shows the ratio of the 90% WC to the best-fit SEL rates for the current data set and for simulated data with added cross section measurements near threshold to better determine LET$_0$, at high LET to better define saturation, and both.

DSEE rate estimation usually suffers from limited statistics. Moreover, for truly destructive SEE, every event represents a failed part, raising the question of how to disentangle part-to-part susceptibility variation from Poisson fluctuations in failure fluence. Even if failure can be avoided and statistics gathered for each part, there is still the question of whether stresses from the destructive SEE mode cause latent damage to the component and alter its susceptibility. Whether we construct the DSEE cross sections with one or many parts, treating each event individually allows us to discover deviations from the assumed model (constant failure rate). In this case, the number of events is always 1, and the variable is the failure fluence. Because SEE are Poisson, failure fluence will be distributed exponentially about the mean failure fluence, $\langle F \rangle = 1/\langle \sigma \rangle$. For each LET, we can accumulate statistics (with a single device if possible or multiple devices if failures occur), and we can fit σ vs. LET using a more complicated GLM:

$$\mathcal{L} = \prod_{j=1}^{\ell} \prod_{i=1}^{k} \langle \sigma(LET_j) \rangle \exp(-\langle \sigma(LET_j) \rangle \times F_{ij})$$
(4)

The product over j is over LET, with $\langle \sigma(LET_j) \rangle = 1/\langle F_j \rangle$, the inverse of the mean failure fluence at LET_j, and the product over i represents statistics accumulated at LET_j

(F_{ij} is the failure fluence for the ith run at LET_j). The model being parameterized enters into the GLM through $\langle \sigma(LET_i) \rangle = 1/\langle F(LET_i) \rangle$. For example, for the usual Weibull form for $\langle \sigma(LET) \rangle$, we would determine the LET_0, σ_{lim} and Weibull parameters that maximize $\mathscr{L}$.

The failure fluences $\{F_{ij}\}$ for LET_j should be distributed exponentially about the mean $\langle F_j \rangle$. We can test whether this is true in a number of ways. Because the Weibull distribution reduces to an exponential distribution when its shape parameter is 1, one can fit $\{F_{ij}\}$ to a Weibull and determine the best-fit value for the shape parameter s. (Note: Do not confuse the Weibull distribution fit to test exponentiality of $\{F_{ij}\}$ with the Weibull cross section form. The latter involves multiple LET values, while the former uses only one. They will yield different best fit parameters.) If s>1 gives a significantly better fit than s=1, it may indicate the failure rate rises with increasing fluence (evidence for latent damage, TID effects or multiple-particle effects). If s<1 yields a better fit, it may indicate a significant component of early failures and high part-to-part variability. Figs. 3a and 3b show fluence to failure for 15 samples of two 100 V power MOSFETs.[13]

While the ST-Micro HG0K failure fluences trend exponentially, the Fujitsu 2SK4219 seems to have excess early failures. Are the excesses significant? Weibull fits reveal that the best-fit shape parameter is 1 for the 2SK4219 and 1.1 for the HG0K. Also, the ratio of the standard deviation to the mean is 1.05 for the 2SK4219 and 0.92 for the HG0K—both near 1 as expected for exponentiality. Neither part deviates significantly from exponential behavior, so the failures can be treated as Poisson.

Often when VGS=0, a MOSFET will not fail during irradiation, but will fail during the post-irradiation gate stress (PIGS) test.[14] Under these conditions, the failure fluence cannot be measured. However, if the test fluence is sufficiently low, then only some of the parts will fail the PIGS test. We can use that proportion, and assuming binomial statistics, estimate a confidence interval for the mean failure fluence. Reference [13] found that when irradiated with VGS=0 V and VDS=100 V, 16 of 22 2SK4219 MOSFETs failed after exposure to 1000 395-MeV Xe ions/cm^2. Similarly 11 of 21 HG0K MOSFETs irradiated with VGS=0 V and VDS=100 V failed after irradiation with 10000 305-MeV Kr ions/cm^2. Table I gives the best estimates and 90% confidence intervals (CI) for the mean failure fluence for the 2SK4219 and HG0K MOSFETs when VGS=0. The resulting fluences can then be inserted into equation 4 to determine the best fit for σ vs. LET.

IV. CONVENTIONAL TID MODELS

Whereas SEE hardness assurance usually ignores part-to-part and lot-to-lot variability, they form the heart of statistical models for TID RHA. This is reflected in the fact that TID RHA statistical models have been standardized for decades in MIL-HDBK 814.[15]

TABLE I: ESTIMATED SEGR FAILURE FLUENCE WHEN VGS=0

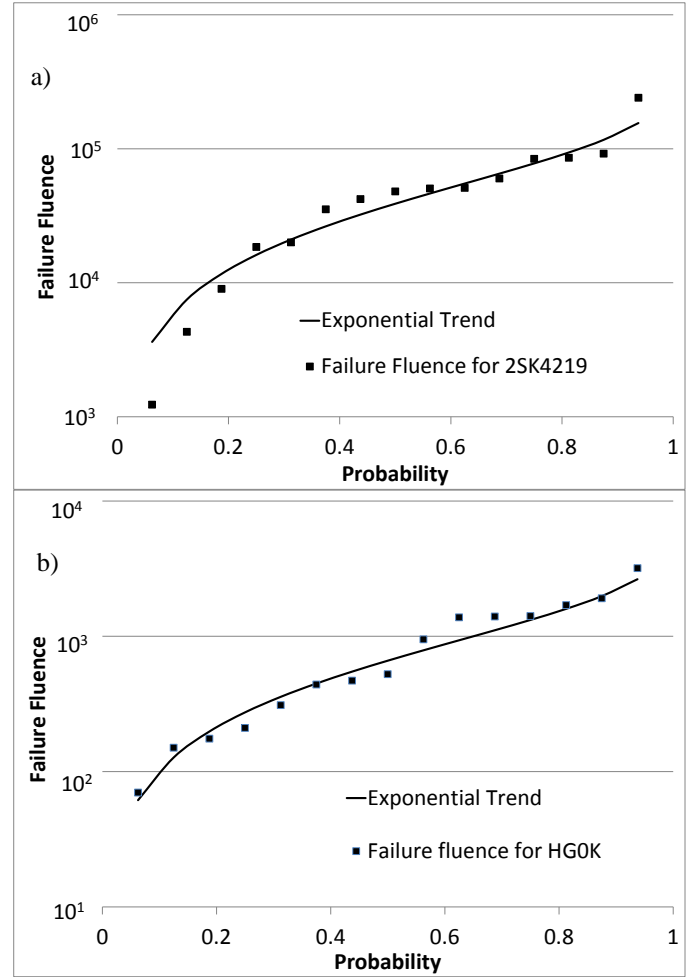| Part #	Best Mean (cm^{-2})	90% CI for Mean (cm^{-2})
2SK4219	640	500-1150
HG0K	13500	8000-21500

Fig. 3 Fluence to failure and best-fit exponential trend for a) Fuji2SK4219 and b) ST Micro HG0K 100 V power MOSFETs irradiated with VDS=100 V and VGS=-10 V. Data courtesy of Veronique Ferlet-Cavrois.[13]

This handbook summarizes over a decade's research into TID RHA methodology, including distribution-free and distribution-dependent sampling, radiation response variability (part-to-part and lot-to-lot), the role of overtest, the role of historical data and many other critical issues. Making no assumptions as to how radiation response varies from part-to-part, large sample sizes would be required to ensure reasonable confidence and probability of success (e.g. 230 samples with no failures to ensure 99% success probability with 90%

confidence). Thus, in practice, most commonly used TID statistical models are predicated on the assumption that radiation response of devices from a single wafer diffusion lot of a part type will follow a well behaved, unimodal distribution. Although this establishes reasonable success probability P_s and confidence level (CL) with small test samples and works well for most commonly used parts, it introduces possible systematic error if the assumptions about the distribution are violated. Moreover, a small sample test is unlikely to discover such violations or reveal details about the extremes of the radiation response distribution. MIL-HDBK 814 suggests using historical data for the part to illuminate such issues. Figs. 4 and 5 show that such a strategy can effectively identify distribution pathologies.


Fig. 4 Small sample sizes preclude identifying the bimodal structure of radiation response in Analog Devices, Inc (ADI) OP484 op amp using a single lot. However, when lots are combined, the bimodal response becomes apparent, with suggestions that bimodality can occur even within a single wafer lot (lot 2).


Fig. 5 Increased gate-to-source leakage current for New England Semiconductor (NES) 2N5019 JFETs exhibits extreme variability, spanning 3 orders of magnitude. Moreover, the lack of an obvious breakpoint suggests a thick-tailed behavior or presence of outliers rather than bimodality.

However, it does not quantitatively treat these features, which may be difficult if they are not obvious. Next, we discuss Bayesian approaches using historical and similarity data and quantitative methods for assessing distribution pathologies (e.g. bimodality and thick tails).

V. DATA STRUCTURE AND BOUNDING RHA

The key to understanding our use of unconventional data types is seen in fig. 6. Even as datasets become less representative, they contain the more representative datasets as subsets. Thus, when we use less representative data to bound flight part performance, what we are really doing is saying that as long as flight parts are not out of family (that is, worse than e.g. 99% of parts in the lot/historical database/process), they should perform no worse than the 99% worst-case parts in that class. The resulting bound is likely conservative, and if it is too loose to ensure mission success, more data or testing are needed. As mentioned above, flight-lot, historical and similarity data are really answering different questions. With flight-lot data, we can ask, e.g., how poorly the flight part can perform if it is no worse than 99% of the parts from our flight-lot (P_s=99%) and the test sample is no more unusual than 90% of similar samples (90% CL). Using historical data when we lack flight-lot data, we can ask how flight parts would perform for flight-lots no worse than, say, the 90% WC lot (again, given the historical data sample is within a desired CL). Finally, for similarity data, we can look at the worst-case part in the worst-case lot for the worst-case part type, all for selected P_s and CL.


Fig. 6 In the absence of the most representative data, flight-part radiation response can still be constrained statistically with less representative data, of which the flight parts are a subset. Just as we use flight-lot data to constrain flight-part behavior, we can use historical data to characterize the distribution of variation across different lots of the flight part type, and data for similar parts to constrain variation across the fabrication process from one part type

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

5

to another. In some cases, other relevant data—trends across feature-size generations, physics, etc.—can be used.

By considering which questions the data can answer, we can expand the sorts of data used to quantitatively bound flight part radiation performance. This means that the following issues must be considered:

1) Is the dataset sufficiently large to infer meaningful bounds on flight-part performance?

Since our approach depends on how radiation response varies from part to part, lot to lot and across part types, the minimum number of parts, lots or part types we can use is 3 (the minimum number for estimating population standard deviation). Although we are expanding the datasets used to bound flight part performance, usually we will not have a large dataset to work with. As such, for sparse data, the method must be conservative. Table II summarizes the minimum data requirements.

TABLE II: MINIMUM DATA SETS FOR DIFFERENT DATA TYPES FOR TID

Data Type	Minimum data	Goal
Flight lot (usually TID)	≥3 parts	Bound flight part performance @ CL
Historical (usually TID)	3 lots (≥3 parts each)	Bound WC part in WC lot for desired CLs for part type
Similarity (SEE or TID)	3 part types (≥3 lots, each with ≥3 parts each)	Bound WC part in WC lot for WC part type for desired CLs for process

2) How do we construct statistics that allow us to maximize the parts in our dataset while still yielding meaningful constraints for flight parts?

Given the limited data usually available for RHA, it is important to include as many similar parts as possible in the dataset, while still ensuring that the distribution of our statistic remains sufficiently compact to draw meaningful conclusions about flight part performance. For example, [16] which looked at TID-induced gain degradation of bipolar junction transistors (BJT), found that normalizing the post-irradiation gain of a transistor to its pre-irradiation gain allowed inclusion of parts with very different pre-rad gains while also reducing part-to-part and lot-to-lot variance. In some cases, similarity data can only provide a partial constraint—for example constraining the onset LET for SEL for components in a process, but not the limiting cross section which depends on the design of the specific parts.[17]

3) How do we ensure that the data used are representative of our flight parts?

Ensuring that flight parts are in family with the parts used in the analysis is challenging even for conventional RHA methods. When using historical or similarity data, the best precaution is to use data for as many lots or part types, respectively, as possible. Seeming outliers may indicate problems with the assumptions of the analysis. Comparing the flight-part datasheet to those for the other parts in the database can also identify significant differences that may invalidate the analysis. In some cases, the vendor may be helpful in validating the similarity of the parts. However, validation of the approach requires flight-part radiation data—either from test or from application success or failure if no test data are forthcoming. In Bayesian SEE and TID risk bounding methods outlined below, analysis of similarity, historical or even lot-specific test data can only influence *a priori* expectations of whether flight parts will succeed or not. As discussed below, a Bayesian approach requires updating prior expectations with flight-part performance data.

VI. BAYESIAN PROBABILITY AND SUBJECTIVITY

To avoid confusion, it is helpful to understand that there are two types of probability. Probabilities inherent to a phenomenon—e.g. when a particular radioactive nucleus will decay—are called objective. No added data can refine our prediction. In contrast, Fig. 7 illustrates that some probabilities are subjective. If we only know that 50% of the balls distributed among the 3 jars in the figure are white, we have no basis for selecting one jar over the others if we want to draw a white ball. However, we could increase our chances of success by measuring the proportion of white balls in each jar. The probability we give for drawing a white ball from a jar depends on our knowledge when we estimate it.

This problem resembles those of bounding flight-part radiation response using historical or similarity data. We know the overall distribution of TID response from historical data, but we do not know where our flight lot falls in that distribution. Or we know the SEE response for a fabrication process based on similarity data, but not where our flight part resides in that distribution.


Fig.7 Subjective or Bayesian probability estimates are based on knowledge at the time of the estimate. In the above example, if we only know that 50% of balls distributed among the three jars are white, our probability of successfully predicting the color of a ball drawn from a jar depend on which jar we pick. We can significantly better our odds by adding knowledge for the individual jars. In this sense, the situation is similar to using RLAT to supplement our knowledge of historical performance of the parts.

Bayes' Theorem plays a central role in updating probability estimates as new data become available. If the probability estimate of event X_i out of a set of possible events $\{X_j\}$ before adding new data $\{d\}$ (the prior probability) is $P_{prior}(X_i)$, then the updated probability P_{post} will be given by Bayes' theorem:

$$P_{post}(X_i \mid \{d\}) = \frac{P(\{d\} \mid X_i)P_{prior}(X_i)}{\sum\limits_{j=1} P(\{d\} \mid X_j)P_{prior}(X_j)} \qquad (5)$$

$P(\{d\}|X_i)$ is the likelihood of the $\{d\}$ if X_i is true, and the denominator is a normalizing factor—the probability of $\{d\}$ regardless of which event in the set $\{X_j\}$ is true. Bayes' Theorem can be generalized for continuous distributions $f(x)$ and data y as follows:

$$f(x \mid y) = \frac{f(y|x)f_{pr}(x)}{\int f(y|x)f_{pr}(x)dx} \qquad (5a)$$

where $f(y|x)$ is the likelihood of y at x and $f_{pr}(x)$ is the prior of $f(x)$.

One of the most controversial aspects of Bayesian probability is the subjectivity of the Prior. There are several ways to limit that subjectivity. If the initial Prior is broad, it has little influence on the left side of (5), and the posterior resembles the likelihood. Also, some versions of Bayesian probability (e.g. empirical Bayesian analysis) allow the Prior to be determined after looking at the data. One locates the Prior near the peak likelihood, and the Prior's width reflects the confidence in the data.

Figure 8 illustrates the relative influences of the data and the Prior. For small datasets, the Prior can significantly influence the posterior distribution, and the latter can change dramatically as new data are added. However, as the data set grows, the distribution stabilizes and the Prior's influence usually becomes negligible. Moreover the flatter (less informative) the Prior, the more rapidly the data dominate. Since RHA almost always deals with small datasets, it is important to minimize the influence of the initial Prior (before adding historical or similarity data)—or at least to ensure that it reflects valid constraints on the data.

One may also ask why bother with the Prior when we could work with the Likelihood directly. There are two reasons to favor a Bayesian approach over a likelihood-based analysis. The first is that Bayesian probabilities follow the laws of probability (e.g. normalization, additivity, etc.), and as such are more intuitive than the corresponding likelihood. The second reason is that a Bayesian analysis is very flexible, and as such can use a broad range of data that would be difficult to exploit using likelihood. Moreover, the logical structure of the data in Fig. 6 favors a Bayesian approach as discussed in

section V. Below, we begin by considering a Bayesian approach for SEE, since the negligible part-to-part and lot-to-lot variation significantly simplify the analysis. In this case, we attempt to bound flight-part SEE rates or consequences (e.g. number of upset bits in an MBU, transient duration or amplitude, etc.) by looking at the SEE response of similar parts fabricated in the same process. Even when it is impractical to bound SEE rates, one may still place some bounds on SEE response.
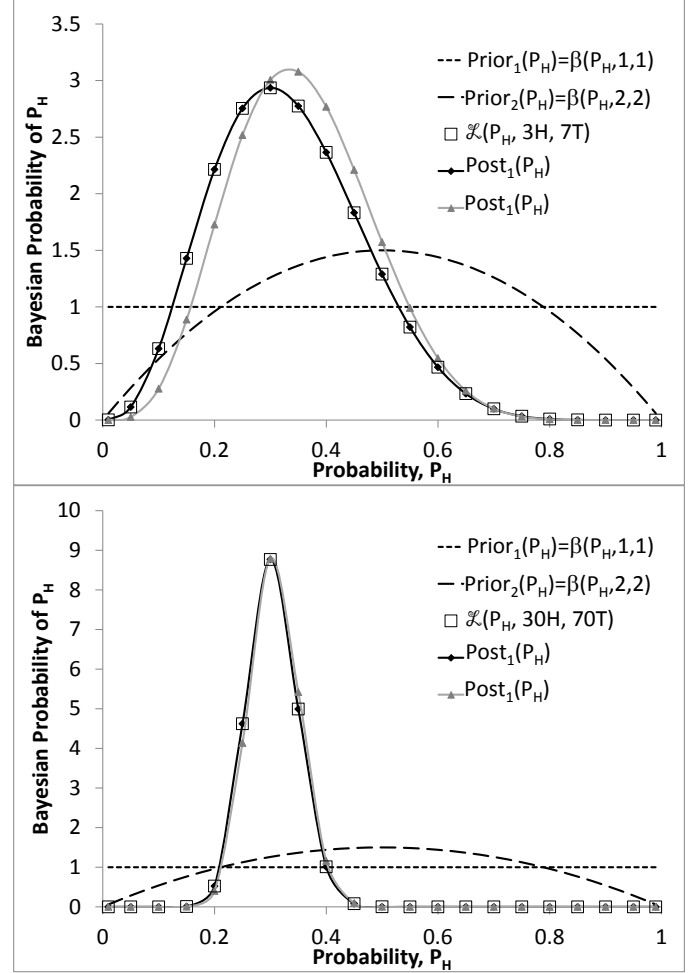

Fig. 8 The influence of the Prior and the likelihood, $\mathscr{L}$ can be seen by considering a coin toss trial. Prior$_1$ considers all probabilities of tossing Heads, P_H equally likely, while Prior$_2$ slightly favors an "honest" coin with $P_H=0.5$. After 10 trials, we have 3 heads and 7 tails, yielding a (normalized) likelihood (open squares) in a). The Posterior yielded assuming Prior$_1$ coincides with the normalized likelihood, while that resulting from Prior$_2$ reflects its influence. In b), increasing the number of trials by a factor of 10, with the same proportion of Heads, we see that the Posterior distributions nearly coincide, regardless of the Prior assumes. (In the legends, the βs refer to the beta distribution, which is the conjugate prior for the probability in the binomial distribution.)

VII. BAYESIAN BOUNDS ON SEE RISK

Reference 15 applied Bayesian techniques to bound risks due to single-event transients (SET) in operational amplifiers and single-event latchup in analog to digital and digital to analog converters (ADCs and DACs). The SET analysis modeled both rates and durations. SET

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

7

amplitudes were also modeled, but the probability of rail-to-rail transients was sufficiently high that assuming such transients is prudent. Although datasets were small for this work (5 op amps for one vendor and 6 for the other), the technique led to reasonable bounds on SET rates for the two vendors' processes and to reasonable if somewhat conservative bounds on SET durations. Bounding SEL risk proved more challenging. Not only did saturation cross sections vary by over a factor of 30 among the ADCs and DACs, but the onset LETs for different parts fabricated in the same process seemed to exhibit a bimodal distribution. Some parts have onset LETs in the single-digits while others were on the order of 15-20 MeVcm^2/mg. The resulting broad range of SEL rates made it impossible to develop compact distributions from which meaningful rate bounds could be determined. However, we were able to model the lower mode of onset LET to determine the WC onset LET to be expected for ADCs and DACs in the 0.6 micron CMOS process from ADI: Over 90% of data converters in this process would likely have onset SEL LET$_0$ > 2.3 MeVcm^2/mg.

It is instructive to discuss why the Bayesian techniques proved more amenable to analysis of SETs than SEL. First, op amps are generally simpler than ADCs and DACs, and SET behavior in many op amps may be dominated by a few susceptible sites.[18] Moreover, the mechanism behind SETs is similar to the normal operation of the part—it is just the charge that is injected to the susceptible region anomalously. In contrast, SEL is a complex, parasitic bipolar phenomenon that depends not just on the characteristics of the process, but also on the device circuit design and layout. Thus it is not surprising that SEL would be much more variable even for similar parts fabricated within a process. What this means is that it is very difficult to place meaningful bounds on flight-part SEL performance without data specific to the flight parts. One exception may be when data for several similar parts fabricated in a particular CMOS or BiCMOS process all indicate no SEL susceptibility. It is not yet known whether this need for flight-part-type specific data applies to SEGR, SEB and other destructive SEE.

Use of the Priors for SEE is fairly straightforward (see Fig. 9). The Prior gives the probability that a particular probability distribution describes the variation of SEE response of similar parts across a particular fabrication process. As such, one can draw contours that contain a particular cumulative probability such that all distributions outside the contour represent worse performance than those inside the contour. If one then selects the worst-performing distribution inside the contour, that distribution represents the worst-case for

the confidence level corresponding to the cumulative probability within the contour.
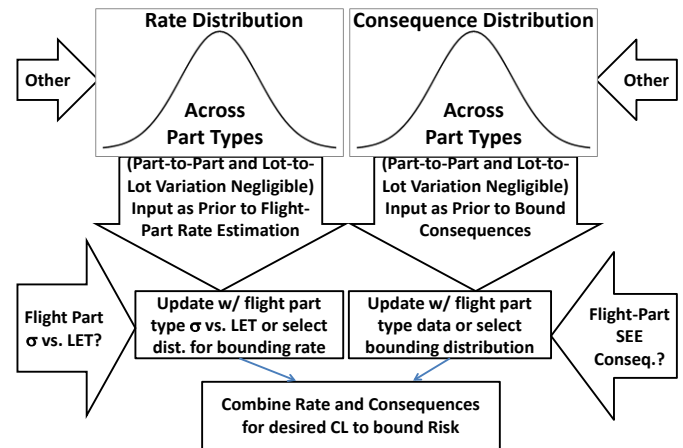

Fig. 9 For SEE, part-to-part and lot-to-lot variation are usually considered negligible, so the Priors are based on the distribution of SEE rates and consequences over similar part types. If data specific to the flight part types become available, these can update the Prior, or if no such data are available, appropriate bounding distributions can be used to bound flight-part risk.

Alternatively, one can average over all possible candidate distributions weighted by their Bayesian probability to produce a best estimate of SEE performance rather than a worst-case bound for a given confidence level. While such averaging is more computationally intensive, it preserves the maximum information originally in the data. As such, it is recommended when one wishes to update the similarity-data Prior with data specific to the flight part types.

The above discussion assumes that part-to-part and lot-to-lot variation in SEE response are negligible. If this is not the case, one must disentangle part-to-part, lot-to-lot and part-type-to-part-type variation, and the approach will resemble that we have developed for Bayesian TID RHA.

VIII. BAYESIAN BOUNDS ON TID DEGRADATION

The need to estimate the different sources of variability significantly complicates the use of historical and similarity data for bounding TID risk. Within a single lot, part radiation response (be it failure dose or parametric degradation at a particular dose) can usually be characterized by a well behaved distribution, with the most probable response occurring near the mean and the range of part responses being characterized by the standard deviation. This does not mean that the distribution is normal, but if we know the mean and standard deviation, we can use the Method of Moments to determine which parameter values for a given distribution family (e.g. normal, lognormal, Weibull, etc.) best reproduce those moments.
Looking at multiple lots of the same part type, part radiation response variability will itself vary across lots. Both the mean and the standard deviation may

vary from lot to lot. This leads to two distributions, one of lot means and the other of lot standard deviations, and we can determine the mean and standard deviation for each of those distributions—four parameters characterizing part-to-part and lot-to-lot variability.

Finally, if we bring in multiple part types (multiple lots of each), each of the four parameters described in the previous paragraph will vary from one part type to another—yielding four distributions, two characterizing the behavior of lot means and two characterizing the lot standard deviations. These distributions describe how these quantities vary across the different part types fabricated in the process according to the data.

Because similarity data are least representative of flight parts, we start with those data to construct our first Priors. The four Priors describe the probability that a particular distribution describes the range of behaviors across part types. The quantity μ_μ, describes the expected (average) lot mean and its distribution describes how it varies from one part type to another for similar part types in the process. The corresponding Prior describes the probability that various distributions describe this variation across part-types. Similarly, σ_μ describes the variation (expressed as standard deviation) from lot to lot of the lot mean, and μ_σ and σ_σ describe the expected lot standard deviation (measuring part-to-part variation in a lot) and the variation of that variability from lot to lot, respectively. We want to select a distribution from the candidates to describe or bound each of the four variables for the parts fabricated in the process (including the flight part type). We can determine these distributions in several ways:

1) We can select the distribution with the highest Bayesian probability. This distribution will best describes the data—especially if we have data for many part types fabricated in the process and the distribution is sharply peaked. This selection is appropriate if we also have data specific to the flight part type and/or flight lot.

2) We can select the distribution that yields worst-case radiation performance consistent with a particular confidence as described in section VII. This is the best strategy when data are limited and we have no data specific to the flight parts or flight part type.

3) Finally, we can average over all distributions weighted by their Bayesian probability. This approach preserves the most information from the similarity data, and it is likely the best strategy when we can update the similarity-based prior with historical and/or flight-lot data.

Regardless of the method used, the result is four probability distributions: $P(\mu_\mu)$, $P(\sigma_\mu)$, $P(\mu_\sigma)$ and $P(\sigma_\sigma)$ (see Fig. 10). If these four variables are uncorrelated, we can construct Priors for lot-to-lot variability:

$$P_{prior}(\mu_\mu, \sigma_\mu) = P(\mu_\mu) \times P(\sigma_\mu) \qquad (6)$$
$$P_{prior}(\mu_\sigma, \sigma_\sigma) = P(\mu_\sigma) \times P(\sigma_\sigma) \qquad (7)$$
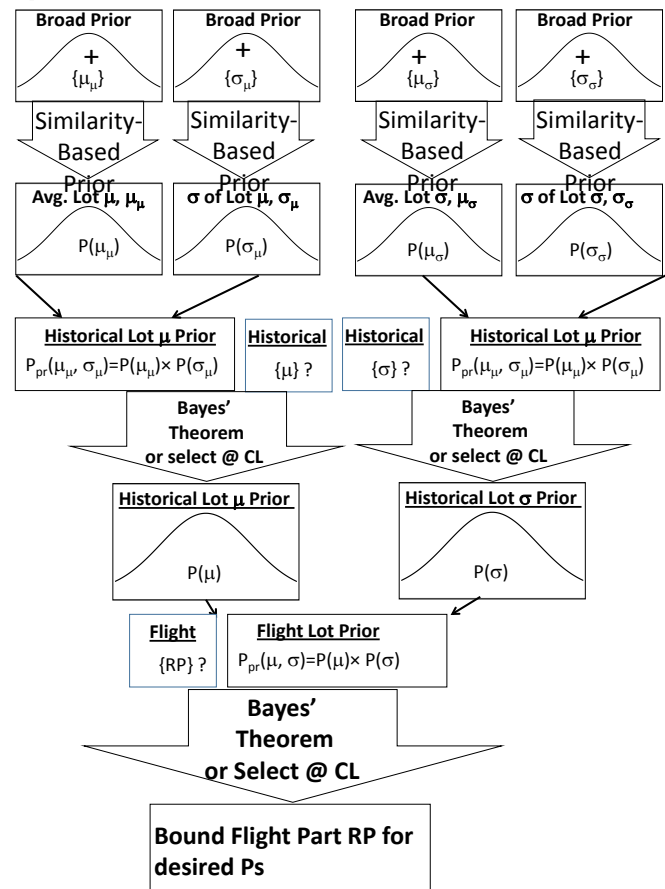

Fig. 10 Bayesian risk assessment for TID is complicated because it must consider part-to-part and lot-to-lot variation as well as how these variabilities change as we look across similar part types fabricated in the same process. Beginning with distributions for expected TID performance and its variation across part types, Priors are updated or used to select Priors for the next stage (historical lot-to-lot variation). The process is then repeated, updating historical Priors with flight-lot data or using them to bound likely flight-part performance based on a "worst-case" lot (for a desired confidence).

These Priors represent the probability that lot-to-lot variation of the lot mean and standard deviation of radiation performance for a part type in the process is described by a distribution with mean μ_μ and standard deviation σ_μ, and similarly for the lot standard deviation. If we have historical data (>3 lots) for the specific flight part type, we can update these Priors using Bayes' Theorem. If not, we can again use one of the three strategies listed previously to select an appropriate distribution to bound the radiation response of lots of the flight part type. Using these $P(\mu)$ and $P(\sigma)$, we construct a Prior as in (6). Then we update the Prior if we have flight-lot data or select a bounding Prior for the flight-lot radiation response. While this approach may seem involved compared to the conventional approach

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

9

(as outlined in MIL-HDBK 814 for instance), it yields much more information—that is, not just the performance of the flight lot, but how this flight lot compares to historical lots and to other similar part types. Moreover, it allows us to bound flight part radiation performance at all stages of the design process, from flight part selection through end of mission life and to tailor the conservatism in the analysis to the risk tolerance of the program.

This approach was applied[16] to estimating gain degradation in BJTs in two families—NPN transistors from Semicoa and PNP transistors fabricated by Microsemi-Lawrence (MS-L). As discussed in section V, we chose as our radiation performance metric the gain degradation factor—the ratio of the pre-irradiation gain to the gain post-irradiation. The analysis found that after 100 krad(Si), gain of an average Semicoa NPN would likely degrade less than a factor of 5 and that after 300 krad(Si) the MS-L PNP gains would likely degrade by less than a factor of 3. Often, these bounds would be sufficient to dispense with radiation testing on the flight parts.

The biggest challenge in carrying out this analysis was finding enough different transistor types that had been tested for sufficiently close conditions that we could make meaningful inferences about variability in gain degradation across lots and across part types. However, despite the limited data, the analysis yielded reasonable and useful bounds on the degradation to expect from a generic Semicoa NPN or MS-L PNP. It also provided context for where a particular lot or transistor type fell within the family of parts. For example, the 2N3700 exhibits slightly less gain degradation than other Semicoa NPNs on average, but has slightly more lot-to-lot variability, whereas for the MS-L PNPs, the 2N2907s exhibited both the greatest average degradation and the most lot-to-lot and part-to-part variation.

IX. UPDATING PRIORS AND HERITAGE DATA

Bayesian probability is an ongoing process—data are always being added to update the Prior and improve our subjective understanding of what we are studying. Ideally, such updates would use the most representative test data. However, such testing may exceed the budgets of low-cost missions. In such cases, mission success or failure serves as the data for updating the Prior. This presumes sufficient insight into failures to determine which parts contributed to the failure. Although such insight is also challenging for low-cost missions, without it, platform reliability cannot improve over time.

Once a mission has ended, its success or failure becomes "heritage data". If a mission succeeds, we can treat it as a suspension or time/dose truncated life test—

that is, as a life or TID test that ended prior to failure. Thus, if the mission of duration t_m were flying n parts, the resulting likelihood of success is

$$\mathcal{L} = \left(\int_0^{t_n} (1 - P_f(t)) dt \right)^n \qquad (8)$$

Here $P_f(t)$ is the failure distribution we are trying to determine. (For TID, we can approximate dose as increasing linearly with time, so we can integrate over time rather than dose.) For such analyses, the constant failure rate for SEE makes updating the Prior with such data easier for these threats. A constant failure rate means the $P_f(t)$ is exponential, and the only parameter required is the mean lifetime. Reference [15] also looks at how to apply heritage missions in different environments, finding that in terms of SEE, a year in geostationary orbit equates to 2.9-4.3 years in a polar orbit (750 km, 98 degrees inclination, sun-synchronous), depending on the geometry of the device sensitive volume (SV), and to 6.5-24.6 years in an International Space Station orbit (500 km, 51.6 degrees inclination), again depending on SV geometry.

Applying heritage data to TID RHA is more problematic. First, mission dose estimates used for radiation design margins (RDM) are upper bounds, leading to a significant overestimate of the component hardness. Even more serious is the fact that failure probability may be negligible up to the end of the mission and then increase dramatically at only slightly higher doses. As such, the worth of TID heritage data is limited unless it encompasses a large number of parts or unless the mission doses for the current mission are much lower than those of the heritage mission. For instance [16] showed that 40 heritage parts would be needed to ensure 99% probability of mission success with 90% confidence if the heritage parts received 2x the mission dose of the current mission. Fortunately, TID testing is usually less expensive than SEE testing and may be possible for many low-budget missions.

X. VERIFYING ASSUMPTIONS

Whenever a statistical model is used for hardness assurance, there is a risk that the part types may not conform to model assumptions. Small samples increase the odds that such deviations may go undetected and introduce systematic errors into the analysis. Thus, verification methods to test consistency with models are essential to successful RHA even for conventional methods. They are even more crucial when using Bayesian methods with diverse types of data. Understanding the sorts of pathologies that occur in a dataset is essential if the data are to be modeled correctly—and once a pathology is discovered in a part

fabricated in a particular process, the other parts in the process should be analyzed for similar pathologies.

As an example, the OP484 data shown in Fig. 4 seem to exhibit bimodality, but how do we show this, particularly when dealing with small test samples? As seen in section IV, one approach is to combine parts across lots into an aggregate distribution.[19] This is justified for the OP484, because the part-to-part variation within a lot is roughly commensurate with the lot-to-lot variation of the lot mean and standard deviation. Fig. 11 shows increased leakage current for 9 lots individually as well as where individual part Ibias values fall relative to the tails of the upper and lower modes of the aggregate distribution.


Fig. 11 The pathology in OP484 radiation response can be characterized by comparing fits of the data to unimodal and bimodal lognormal distributions. This reveals that the lowest Ibias in lot 2 has less than a 1% probability of belonging to the same upper mode as the other parts in the lot, while the parts in lot 3 all most likely fall in the lower mode—albeit just barely.

An important question is whether any of the lots exhibit bimodality, or whether the bimodality is only from lot to lot. Fig. 4 suggests the rough locations and widths of the two modes, and it also shows that lot 2 seems to have parts occurring in both modes. To better determine this, we fit the aggregate data to both a unimodal lognormal and a bimodal lognormal using maximum likelihood. While the bimodal distribution will fit better, how can we compare the two fits, since the bimodal fit has five parameters while the unimodal has two? Akaike's Information Criterion (AIC) corrects for this effect by penalizing the more complex model with a term proportional to the number of parameters, k, while the goodness of fit is measured by the likelihood. We use the corrected form AICc for small data sets, n:

$$\text{AIC}_c = -2\log[\mathcal{L}(\mathbf{\Theta}, \mathbf{X})] + 2k + \frac{2k(k+1)}{n-k-1} \quad (9)$$

where $\mathcal{L}(\mathbf{\Theta}, \mathbf{X})$ is the maximum likelihood for the k-parameter vector $\mathbf{\Theta}$ and the data vector $\mathbf{X}$ and n is the data count in X.

For as few as 3 lots (any variability for 2 lots will likely appear bimodal), AIC strongly favors a bimodal model. The evidence for bimodality continues to increase up to 6 lots and then levels out. AIC selects the model with the greatest predictive power, rather than the model that gives the best fit. This can be seen in Fig. 12, where we have plotted the 99% worst predicted Ibias after 100 krad(Si) for 90% confidence (calculated using one-sided tolerance limits and the best-fit lognormal mean and standard deviation) for the unimodal model and the worst-case mode of the bimodal model.


Fig. 12 Based on AIC, three lots of data for the OP484 are sufficient to show that the bimodal model has better predictive power for change in Ibias. This is also supported by the fact that the bimodal 99/90 WC Ibias remains very stable as new data are added, whereas the corresponding quantity assuming unimodal behavior overpredicts the change and fluctuates wildly.

The bimodal estimate is lower, more reasonable and more stable, further suggesting that the bimodal model is correct. Moreover, the bimodal fit yields not just the position and width of each mode, but also the relative proportion in each mode (about 50% for the OP484). In Fig. 11, we indicate with a solid line the Ibias with equal probability of belonging to either distribution. Data above the dashed line have >5% probability of belonging to the upper mode, and below the dotted line, a >5% probability of belonging to the lower mode. Only lot 3 contains data that fall between these two lines. This suggests that the lowest data point from lot 2 most likely belongs to the lower mode, indicating bimodal response in a single lot is possible for this device.

Often, the nature of the pathology for a dataset may not be as obvious as OP484 bimodality. For the NES 2N5019 JFETs shown in Fig. 5, it is not clear whether the 3 order-of-magnitude spread arises due to bimodality, outliers or "maverick" parts or is indicative of a thick tail for the distribution of increased gate-to-source current after irradiation. To better ascertain the nature of the distribution, we use a rank plot for the data, along with a lognormal fit to the data both with and

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

11

without a candidate outlier (Fig. 13). Omitting the worst data point does not improve the fit of a lognormal to the data (R^2=0.823 with the worst point and 0.818 without it), suggesting that this point is not an outlier, and a distribution with a thicker tail than lognormal is likely required to fit the tail of the distribution.


Fig. 13 A fit to an aggregate distribution of increased IGSS for the 2N5019 JFET shows that the tails of a lognormal distribution are too thin to model the large part-to-part variation. Omitting the worst case part does not improve the agreement between the model and data, indicating that the worst-case part is likely not an outlier and the distribution is actually thick-tailed.

In contrast to the situation for the 2N5019, [19] found evidence of an outlier when comparing increased leakage input bias current after 50 krad(Si) for parts fabricated in LTC's RH process. Among parts in this process, the RH27 exhibits slightly more increased Ibias after 50 krad(Si) (~10% more than the next worst part). However, Fig. 14 shows that part-to-part variability for the RH27 is much greater—to the extent that excluding the RH27 from the sample significantly improves the goodness of fit of the data to a lognormal (R^2 moving from 0.64 to 0.98). A subsequent call to LTC revealed that the RH27 uses the same mask set as the OP27, and so has no circuit-level hardening. Other parts in the process use both circuit- and process-level hardening.

Preliminary analyses such as those applied above for TID can also be used to ensure that similarity analyses for SEE use the largest range of similar parts possible while excluding outliers and identifying pathologies. Moreover, to ensure that SEE are Poisson, analyses similar to those used in section IV to ensure exponentiality in fluences to failure can be used to test this assumption for any part type, as long as the fluences between errors are known. If the distribution of fluences between errors deviates from exponential, it is likely that the error generating process is not purely Poisson.

XI. PRESENTATION OF RESULTS

Statistical inference based on historical or similarity data can be a valuable aid for radiation test planning, on-the-fly test decisions, selection of parts or design strategies and a variety of other mission critical activities. However, radiation analysis results must be conveyed effectively to the designers ultimately responsible for those decisions. Here we discuss tools that facilitate visualization of statistics and understanding of results and their confidence levels. To facilitate accessibility, we have used widely available, open-access tools. Most of the plots in this section are created using the statistical coding language R [20], along with the graphics add-on ggplot2.[21].


Fig. 14 Examination of the change in input bias current (Ibias) of op amps and comparators fabricated in the LTC RH process reveals that the RH27 op amp shows far more part-to-part variation than other parts fabricated in the process—to the extent that including it distorts the lognormal fit for the other parts in the process. This suggests the RH27 is an outlier.

A. SEE results

Because estimated SEE rates typically depend on the parameters of a cumulative Weibull fit to the σ vs. LET data, it is important to convey as much information as possible in plots of the data so that analysts and engineers unfamiliar with the testing can assess data reliability, the conservatism of the fit and so on. For example, Fig. 2 in section II presents not just SEL σ vs. LET data for the LTC1419 ADC, but also error bars and best and 90% WC fits. The table below the figure shows how the limited data affect confidence in the error rate. These data are sufficient to give designers an idea of the range of probabilities of an SEL during the mission.

In contrast, when dealing with SET, the rate is sufficiently high that the occurrence of the event is a virtual certainty during the mission. In such cases, designers need to understand transient characteristics so that appropriate filtering or other mitigation can be implemented. In the past, this was often conveyed by specifying a transient that was worst-case in terms of both amplitude and duration, whether such a transient

were realistic or not. Statistical plots allow us to convey much more information to designers, allowing them the option to design for more realistic transients and to tailor filtering to the risk tolerance of the application. Fig. 15 shows a duration vs. amplitude scatterplot of the raw transient data gathered during a laser SEE test of the Texas Instruments LMP2012 op amp, the resulting histograms of both amplitude and duration and a contour plot of the probability density function (PDF) derived from the raw data using the "density()" tool built into the R core software. The PDF plot is particularly useful for tailoring mitigation to risk tolerance. The histograms are useful for determining possible outliers and extreme events.


Fig. 15 SET data for the Texas Instruments LMP2012 op amp observed during laser SEE testing. Data in the duration-amplitude scatter plot are projected to also give duration and amplitude histograms.

B. TID results

As seen in section IX, viewing lot test data in the context of archival data for historical lots and similar part types can lead to a different conclusion than one would reach based solely on the numerical results of the test. Conveying this context can be crucial to building consensus for difficult RHA decisions. If a part exhibits significant lot-to-lot TID variation and one is working with small test samples, context becomes even more important, and usually the best way to convey the context is graphically. Many types of graphs and charts are used to convey statistical information. The scatter plots, histograms and PDF contours discussed in the context of SETs illustrate some options. In sections III and IX, we used rank or quantile plots to assess, respectively, whether failure fluences followed expected exponential behavior and to look for outliers or multiple modes. These plots offer a graphical method of comparing two probability distributions or comparing

data to a theoretical distribution to gain insight into the data's behavior. We can also compare data for two different historical lots or two different part types to assess whether inclusion in a historical or similarity data analysis is appropriate (e.g. to determine Priors for statistical inference).

To illustrate how historical context can affect RHA decisions, consider an issue encountered for the Magnetospheric Multi-Scale (MMS) Mission. Lot #1016 of Microsemi JANTXV2N3700 NPN BJT failed to meet its specified gain hfe2 (collector to emitter voltage, VCE=10 V and collector to emitter current, ICE=0.1 mA) after 25 krad(Si). Although the parts have no radiation guarantee, the rapid failure raised questions about the trustworthiness of the lot even for low-gain applications. Radiation analysts compared lot 1016's behavior to those of historical lots (0702 and 2462). Fig. 16 shows the cumulative distribution (CDF defined as rank divided by # samples) for all three lots at 25, 50, 75 and 100 krad(Si). This shows that lot 1016's performance is similar to that of lot 2462 and that there is large lot to lot variation for all doses.


Fig. 16 Cumulative distribution function (CDF) of gain hfe2 for three lots of data at four dose steps.

In Fig. 17 we compare lot-to-lot variation more directly, plotting (a so-called quantile-quantile or Q-Q plot) the gain at a given quantile for one lot against the gain for the same quantile for another lot (for each dose). The greater the curvature seen in the series, the more dissimilar the behavior of the lots. These plots show that hfe2 for lot 1016 degrades at lower doses than for the other two lots, and that the greatest difference is seen at 50 krad(Si). At 75 krad(Si) and 100 krad(Si), the other lots start to catch up, but still remain less degraded than lot 1016. Again, lots 1016 and 2462 have similar

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

13

performance, while lot 0702 is much harder. The difference is most pronounced at 25 and 50 krad(Si).

Fig. 18 shows the combined probability density function of each lot's mean and standard deviation of hfe2 normalized to its pre-rad value. The mean gain exhibits very different behavior with dose for each lot as indicated by the separation of the probability density into two groupings as dose increases in Fig. 18.
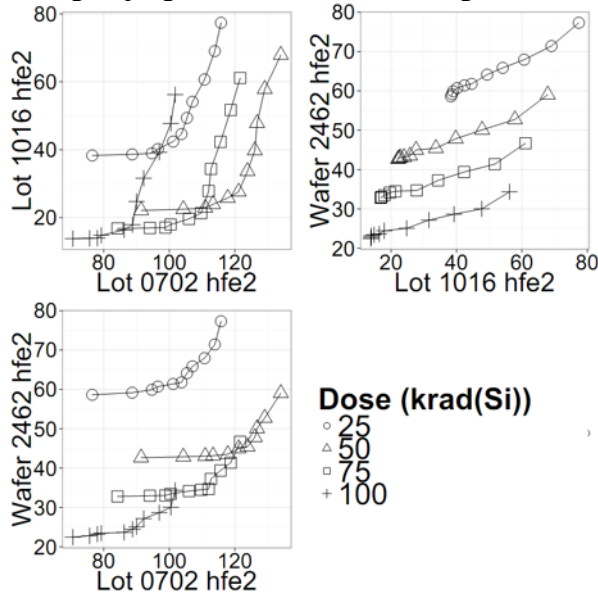

Fig. 17 Quantile – Quantile plot comparing gain for pairs of lots selected from the three lots. Straighter lines in the plot indicate similar performance of the two lots, while curved lines indicate different performance.

The PDF used in both Fig. 18 (as well as that in Fig. 15) results from combining the kernel density estimation of the dataset for each mean and standard deviation. We used the "density()" tool built into the R core software. For a given datapoint the kernel density is approximated with an optimized bandwidth from a Gaussian distribution. This information could be used to establish our Priors for Bayesian analysis.

XII. CONCLUSIONS

Although the basic statistics of RHA have been understood for decades, techniques for analyzing both SEE and TID have advanced significantly. These techniques facilitate not only improvements for conventional RHA methods, they also introduce the possibility of using unconventional data sets to quantitatively bound flight-component radiation performance. In terms of conventional RHA, the techniques we have applied above are both generalizations of likelihood analysis: Generalized Linear Models use likelihood to parameterize complicated models, such as the Weibull form for σ vs. LET in SEE analysis. The Akaike Information Criterion (and related quantities) makes it possible to compare the performance of models with different complexities by penalizing complex models (which usually give better fits to the data) relative to simple ones. This makes the criterion for selecting a model not just its goodness of fit, but rather its predictive power. While the example we cite above used AIC to distinguish between unimodal and bimodal TID degradation in the OP484 op amp, the technique could also prove useful in SEE analysis challenges such as identifying whether multiple mechanisms contribute to a particular SEE, determining the optimum number of charge collection volumes in a complex Monte Carlo based rate estimation and so on.

We have also discussed use of Bayesian analysis for both TID and SEE hardness assurance. Bayesian methods are well suited to many problems that arise in RHA because many of the probabilities encountered are subjective—that is, the probabilities can change if we add information to our current understanding.

There are several advantages to Bayesian methods: They bound risk at all stages of the analysis, so that the radiation analyst can determine if more information (e.g. testing, analysis, etc.) is needed. The methods assign candidate models a probability rather than a likelihood, making interpretation easier. Also, Bayesian methods are sufficiently flexible that almost any relevant data can be used. Moreover, as long as the flight parts are not out of family compared to the data we bring to the analysis, the bound we estimate will likely be conservative, since the flight parts are always a subset of the data we use. This is important for RHA, because often the datasets available are not large.


Fig. 18 A combined PDF for three lots of 2N3700 transistors from Microsemi generated using the density() tool in the R statistics package.

The key to successful application of these new techniques lies in understanding the questions the data

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

14

under consideration can answer and how those answers constrain the likely radiation performance of the flight parts under consideration. For example, historical data for the flight part type can tell us how a 99% worst-case part will perform for a 90% worst-case lot of the flight part, and similarity data can yield such data for a 90% worst case part type fabricated in the same process. The resulting bounds may suggest that we have high confidence of the flight parts fulfilling their mission requirements—or if not, the Priors on which these bounds were based can be updated with data specific to the flight parts.

Statistical models like those discussed above offer the ability to bound flight-part performance economically and reliably. Also, because the method relies on historical test data, it ensures that this valuable resource is put to maximum use and provides added incentive to standardize test methods and make the data as easy to interpret as possible. Moreover, because one can adjust the desired confidence level and degree of conservatism, the method can be tailored to the risk tolerance of the hardware being analyzed. Finally, because the methods are extremely flexible, they can be adapted to include increasing levels of complexity. For example, analytical models of SET performance can be used to interpolate performance of components to different application conditions. Priors can be developed based on technology trends, and so on.

However, perhaps the most significant contribution these methods make is to emphasize the continual nature of RHA. This is especially true for Bayesian methods. While we use the Prior to infer likely radiation performance of flight parts, these predictions must be validated and the Priors updated by the infusion of new data—be it test data or heritage data for the mission in question. This ensures that both the methodology and the spacecraft reliability improve over time.

REFERENCES

[1] R.D. Morris and C.C. Foster, Cross-section estimation in the presence of uncertain dosimetry, IEEE Transactions on Nuclear Science, vol. 57, no. 6, pp. 3528-3536, Dec. 2010

[2] JESD57, "Test Procedures for the Measurement of Single-Event Effects in Semiconductor Devices from Heavy Ion Irradiation," Electronic Industries Association, Engineering Department, Arlington, VA, 1996.

[3] E. Petersen, "Single-Event Analysis and Prediction," Snowmass, CO, 1997 IEEE NSREC Short Course.

[4] D. M. Fleetwood and H. A. Eisen, "Total-dose radiation hardness assurance," IEEE Trans. Nucl. Sci. 50, pp. 552-564 (2003).

[5] J. R. Schwank, M. R. Shaneyfelt and P. E. Dodd, "Radiation Hardness Assurance Testing of Microelectronic Devices and Integrated Circuits: Test Guideline for Proton and Heavy Ion Single-Event Effects," IEEE Trans. Nucl. Sci. 60, pp. 2101-2118 (2013).

[6] J. H. McDonald, "G-test for goodness-of-fit" in *Handbook of Biological Statistics* (2nd ed.). Sparky House Publishing, Baltimore, Maryland. (2009) (pp. 46–51).

[7] R. Ladbury, "Statistical properties of SEE rate calculation in the limit of large and small event counts," IEEE Trans. Nucl. Sci., vol. 54, no.6, pp. 2113–2119, Dec. 2007.

[8] N. Sukhaseum et al. "Statistical Estimation of Uncertainty for Single Event Effect Rate in OMERE." 13th European Conference on Radiation and Its Effects on Components and Systems, RADECS. 2011.

[9] J. L. Titus et al., "Prediction of early lethal SEGR failures of VDMOSFETs for commercial space systems," IEEE Trans. Nucl. Sci., vol. 46, no. 6, pp. 1640–1651, Dec. 1999.

[10] S. Kuboyama et al., "Averaged LET Spectra for SEE Rate Prediction for Power Devices," IEEE Trans. Nucl. Sci., vol. 56, no. 6, pp. 2056 - 2060, Dec. 2009.

[11] J.-M. Lauenstein, "Effects of Ion Atomic Number on Single-Event Gate Rupture (SEGR) Susceptibility of Power MOSFETs," IEEE Trans. Nucl. Sci., vol. 58, no. 6, pp. 2628 - 2636, Dec. 2011.

[12] L. D. Edmonds, L. Z. Scheick and M. W. Banker, "Single Event Rates for Devices Sensitive to Particle Energy," IEEE Trans. Nucl. Sci., vol. 59, no. 6, pp. 2936 - 2944, 2012.

[13] V. Ferlet-Cavrois et al., "Statistical Analysis of Heavy-Ion Induced Gate Rupture in Power MOSFETs—Methodology for Radiation Hardness Assurance," IEEE Trans. Nucl. Sci., vol. 59, no. 6, pp. 2920 - 2929, Dec. 2012.

[14] A. D. Privat, et al., "On the Use of Post-Irradiation-Gate-Stress Results to Refine Sensitive Operating Area Determination," IEEE Trans. Nucl. Sci., vol. 61 no.6, 2930-2935, Dec. 2014.

[15] MIL-HDBK 814, "Ionizing Dose and Neutron Hardness Assurance Guidelines for Microcircuits and Semiconductor Devices," 08 February 1994.

[16] R. Ladbury and B. Triggs, "A Bayesian Approach for Total Ionizing Dose Hardness Assurance," IEEE Trans. Nucl. Sci., vol. 58, no. 6, pp. 3004 - 3010, Dec. 2011.

[17] R. Ladbury and M. J. Campola, "Bayesian Methods for Bounding Single-Event Related Risk in Low-Cost Satellite Missions," IEEE Trans. Nucl. Sci., vol. 60, no. 6, pp. 4464 - 4469, 2013.

[18] S. Buchner and D. McMorrow, "Single-Event Transients in Bipolar Linear Integrated Circuits," IEEE Trans. Nucl. Sci., vol. 53, no. 6, pp. 3079 - 3102, Dec. 2006.

[19] R. Ladbury and J. L. Gorelick, "Statistical Methods for Large Flight Lots and Ultra-High Reliability Applications," IEEE Trans. Nucl. Sci., vol. 52, no. 6, pp. 2630 - 2637, Dec. 2005.

[20] R. Ladbury, J. L. Gorelick and S. S. McClure, "Statistical Model Selection for TID Hardness Assurance," IEEE Trans. Nucl. Sci., vol. 57, no. 6, pp. 3354 - 3360, Dec. 2010.

[21] R Core Team (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: http://www.R-project.org/.

[22] H. Wickham. ggplot2: elegant graphics for data analysis. Springer New York, 2009.

To be published in the Institute of Electrical and Electronics Engineers (IEEE) Transactions on Nuclear Science, Special Edition, 2015.

15